

5

10

15

Patent Application of

Haw-Minn Lu

20

Fanout Upgrade for a Scalable Switching Network

25

Background - Field of Invention

This invention relates to redundant multistage switching networks, specifically to the non-stop addition of new ports to all routers of each row of such a network.

30

Background - Discussion of Prior Art

Many multistage switching networks are defined primarily as binary based. Some of these networks have been generalized for arbitrary fanouts.

35

Fig. 1A shows an 8-port binary Banyan network. Banyan networks are an example of multistage switching networks that have been generalized to arbitrary fanouts. Fig. 1B shows a 27-port ternary Banyan network. Both examples have the same number of rows.

40

There is no known investigation of the process of upgrading the fanout of Banyan networks. There are two likely reasons why such an upgrade path is undesirable. First, most implementations of switching networks using the Banyan architecture employ specific binary sorting algorithms to route traffic. Second, an upgrade of an n -stage 2^{n+1} -port binary network to an n -stage 2×3^n -port ternary network or n -stage 2×4^n -port quaternary network would entail an exponential growth in number of ports required. From the examples already shown, this is a very large number even for the small numbers considered here.

Objects and Advantages

Accordingly, the several objects and advantages of my invention are:

- 5
- (a) to provide a procedure by which a redundant multistage switching network can be upgraded by expanding the number of ports per router without having to break an external connection; and
 - 10 (b) to provide a procedure by which a redundant multistage switching network can be upgraded by expanding the number of ports per router with a minimum loss in throughput bandwidth.

Further objects and advantages of our invention will become apparent from a consideration of the drawings and ensuing description.

15

20075036 024002
200720 33052007

Description of Drawings

Fig. 1A shows a 16 port binary Banyan multistage switching network with three rows. (prior art)

5 Fig. 1B shows a 54 port ternary Banyan multistage switching network with three rows. (prior art)

Fig. 2 shows a 30 port redundant blocking compensated cyclic group (RBCCG) multistage switching network. Hashed areas represent slots for adding ports (prior art)

10 Fig. 3A shows a 30 port RBCCG multistage switching network where all routers have been augmented by 2 ports each, resulting in a 40 port switching network.

Fig. 3B shows a section of a 40-port RBCCG multistage switching network with the proper interconnection network between two rows of routers.

15 Fig. 4A, Fig. 4B, Fig. 4C, Fig. 4D, Fig. 4E, Fig. 4F, Fig. 4G, Fig. 4H, Fig. 4I, Fig. 4J, Fig. 4K, Fig. 4L, Fig. 4M, Fig. 4N, Fig. 4O, and Fig. 4P show the connections between router rows $R(1,*)$ and $R(2,*)$ being rewired into a 40 port RBCCG multistage switching network interconnection by adding a connection between bottom port 3 of $R(1,4)$ and top port 3 of $R(2,4)$; adding a connection between bottom port 3 of $R(1,3)$ and top port 3 of $R(2,0)$; moving the connection from top port 2 of $R(2,4)$ to top port 3 of $R(2,3)$; moving the connection from top port 2 of $R(2,3)$ to top port 3 of $R(2,2)$; moving the connection from top port 2 of $R(2,0)$ to top port 2 of $R(2,3)$; moving the connection from top port 2 of $R(2,2)$ to top port 3 of $R(2,1)$; moving the connection from top port 2 of $R(2,1)$ to top port 2 of $R(2,4)$; adding a connection between bottom port 3 of $R(1,2)$ and top port 2 of $R(2,1)$; moving the connection from top port 1 of $R(2,4)$ to top port 2 of $R(2,2)$; moving the connection from top port 1 of $R(2,2)$ to top port 1 of $R(2,4)$; adding a connection between top port 1 of $R(2,2)$ and bottom port 3 of $R(1,1)$; moving the connection from bottom port 0 of $R(1,1)$ to bottom port 3 of $R(1,0)$; moving the connection from top port 1 of $R(2,3)$ to top port 2 of $R(2,0)$; moving the connection from top port 1 of $R(2,1)$ to top port 1 of $R(2,3)$; moving the connection from top port 1 of $R(2,0)$ to top port 1 of $R(2,1)$; moving the connection from bottom port 1 of $R(1,1)$ to bottom port 0 of $R(1,1)$, respectively.

20 Fig. 4Q shows a partially upgrade 40 port multistage switching network with a completed rewire of the interconnections between router rows $R(1,*)$ and $R(2,*)$ adding a connection between bottom port 1 of $R(1,1)$ and top port 1 of $R(2,0)$.

25 Fig. 5B shows a partially upgrade 40 port multistage switching network with a completed rewire of the interconnections between router rows $R(0,*)$ and $R(1,*)$

30 Fig. 5C shows a partially upgrade 40 port multistage switching network with a completed rewire of the interconnections between router rows $R(2,*)$ and $R(3,*)$

35 Fig. 6A, Fig. 6B, Fig. 6C, Fig. 6D, Fig. 6E, Fig. 6F, Fig. 6G, Fig. 6H, Fig. 6I, Fig. 6J, Fig. 6K, Fig. 6L, Fig. 6M, Fig. 6N, Fig. 6O, Fig. 6P, Fig. 6Q, Fig. 6R, Fig. 6S, Fig. 6T, Fig. 6U, Fig. 6V and Fig. 6W show the connections between router rows $R(1,*)$ and $R(2,*)$ being rewired in an alternate method into a 40 port RBCCG multistage switching network interconnection by an alternate method. First, by swapping bottom ports 1 and 0 of $R(1,1)$; swapping bottom ports 2 and 1 of $R(1,1)$; swapping bottom ports 0 and 3 of $R(1,2)$; swapping bottom ports 2 and 0 of

R(1,2); swapping bottom ports 0 and 2 of R(1,3); swapping bottom ports 1 and 3 of R(1,3); swapping bottom ports 0 and 1 of R(1,4); swapping bottom ports 0 and 2 of R(1,4); swapping bottom ports 0 and 3 of R(1,4); swapping top ports 2 and 3 of R(2,4); swapping top ports 1 and 2 of R(2,4); swapping top ports 2 and 3 of R(2,3); swapping top ports 2 and 3 of R(2,2); swapping top ports 1 and 2 of R(2,1); and finally, swapping top ports 2 and 3 of R(2,0), respectively. Second by adding a connection between bottom port 0 of R(1,4) and top port 3 of R(2,1); adding a connection between top port 2 of R(2,3) and bottom port 1 of R(1,3); adding a connection between bottom port 2 of R(1,2) and top port 2 of R(2,0); moving the connection from top port 1 of R(2,1) to top port 2 of R(2,2); moving the connection from top port 1 of R(2,2) to top port 1 of R(2,4); adding a connection between top port 1 of R(2,2) and bottom port 3 of R(1,1); moving the connection from top port 0 of R(2,3) to top port 1 of R(2,1); adding a connection between top port 0 of R(2,3) and bottom port 3 of R(1,0), respectively.

Fig. 7A, Fig. 7B and Fig. 7C show the algorithm for expanding the fanout of the component routers of a compensated cyclic group multistage switching network, Fig. 7A shows the most general algorithm, while Fig. 7B shows a more specific algorithm.

Fig. 8A and Fig 8B show two different algorithms for selecting the order the interconnections should be rewired.

Fig. 9A, Fig. 9B, and Fig. 9C show three different algorithms for selecting which port to rewire during the rewiring phase.

Fig. 10 shows an algorithm for relabeling ports to improve the efficiency of the upgrade process.

Fig. 11 shows the flowchart for the fanout upgrade algorithm shown in Fig. 7A and Fig. 7C.

Fig. 12 shows the flowchart for the fanout upgrade algorithm shown in Fig. 7B and Fig. 7C.

Fig. 13A and Fig. 13B show the flowcharts for the row selection algorithms shown in Fig. 8A and Fig. 8B respectively.

Fig. 14A, Fig. 14B and Fig. 14C show the flowcharts for the port selection algorithms shown in Fig. 9A, Fig. 9B and Fig. 9C, respectively.

Fig. 15 shows flowchart for the relabeling algorithm shown in Fig. 10.

Summary

A multistage switching network can be expanded by adding ports to existing switching elements without breaking any external connections.

If the switching elements are routers or have the ability to route traffic the expansion may be performed without severely disrupting its ability to route traffic, provided that the routers themselves have the capability of having ports added while still in operation. Any number of ports may be added in any position as long as there are an equal number added to both top and bottom ports. In general, these ports must be reassigned logically so the position of the new ports is irrelevant.

The fanout upgrade process has three constituent components. The first component is to install new line cards to provide each router with new ports.

The second component has two types steps of which one is optional. The first type is to select a port that is misconnected in reference to the desired topology and to rewire it to its appropriate counterpart. The second type is optional and is the permutation of ports on the same router which results in a correct port being connected to the correct router. If such a permutation can be performed logically, traffic need not be disrupted. These two types of steps are repeated in any order until the desired topology is achieved. Since the ordering of such steps can be arbitrary, more structure can be added by breaking this component into three subcomponents.

The first subcomponent is the selection of the order in which each interconnection network is rewired into the desired topology. The order can be arbitrary, however, since the path redundancy is greatest in the middle of the network, the best order has been found to start with the middle row or the closest to the middle. The subsequent rows should be ordered with those closest to the middle first and those farthest away from the middle last.

The second subcomponent, which is optional, is the permutation of ports described above. This step should always be taken before rewiring an interconnection network, if such a permutation can be performed logically so that no traffic disruption occurs. This step minimizes bandwidth reduction during the rewiring of an interconnection network.

The third subcomponent is the rewiring phase. This is accomplished by systematically disconnecting from a selected port any misconnected connection to that port and to the corresponding port to which the selected port is to be connected. With any misconnections disconnected, the selected port and its appropriate corresponding port may then be connected. In order to minimize the impact on the throughput bandwidth, priority in the port selection process should be given to selecting ports not currently connected. This introduces a plethora of port selection algorithms. Additional preferences, contribute to better throughput bandwidth during the upgrade process.

The final component is relatively simple. All external connections that are not already connection should be connected and placed into service.

Description of Invention

A redundant blocking compensated cyclic group (RBCCG) multistage network, is shown in Fig. 2 and discussed further in United States Patent 5,841,775, "Scalable Switching Networks" by Alan Huang, Nov. 24, 1998. It consists of rows of routers or other switching elements. These rows of routers are connected together via interconnection networks. The routers are designated R(row, column) where the top left most router is denoted R(0,0). The top ports of each router are numbered from left to right starting with 0. The bottom ports of each router are numbered from left to right starting with 0. The top ports of each interconnection network are numbered from left to right starting with 0. The bottom ports of each interconnection network are numbered from left to right starting with 0.

This method applies routers or switching elements for which ports may added without having to shutdown the router. In Fig. 2, the hatched areas on each router indicate potential insertion points for new ports. Typically, routers are expanded by the addition of new line cards. Insertion of a new line card may add more than a single port. In general, many ports can be added using the methods described here. In addition, a single line card with two or more ports may be installed to provide both additional top and bottom ports to a router. For the purposes here, the method is applied to the expansion of a single top port and a single bottom port for each router, as shown in Fig. 3A. With the new ports added, interconnection networks 10, 11, and 12 need to be rewired to look like interconnection network 20 as shown in Fig. 3B.

The most general form of the upgrade process is given by Fig. 7A, with supporting functions given in Fig. 7C, and flowcharted in Fig. 11. The first major step is to add the desired new ports to each routers. This is depicted in Fig. 7A and Fig. 11 with a algorithmic loop. Though no specific order is given for which the routers are augmented, they may be augmented in any order. No particular order has any advantage in terms of minimizing traffic disruption or number of steps. The choice may come down to which is physically more convenient.

The upgrade process continues after the ports are added by selecting any port which is misconnected with respect to the desired final topology. The order of selection does have an impact on the disruption of traffic and even the number of total steps. A more systematic and organized approach is given below in a refinement of this algorithm. Once a port is selected it is rewired by connecting it to its appropriate port with respect to the final topology. Henceforth, this port will be known as its corresponding port. The remainder of the algorithm in Fig. 7A and flowchart Fig. 11 is a detailed description of the rewiring based on the connection states of the port and its corresponding port.

There is also an optional branch in the upgrade process. In scanning the routers, one may discover routers for which some connections originate from the correct routers (with respect to the desired final topology) but connect to the wrong port. If a logical reassignment of port addresses were possible, achieving a virtual port permutation while leaving the physical connections unchanged. This would reduce the number of connections that need to be broken. This step, as indicated by the optional branch, may be executed at any time, that is it need not be executed in every loop of the algorithm. For instance, the sequence could be rewire, relabel, relabel, relabel, rewire, rewire, relabel, rewire, etc. A more systematic approach is given in a refined version of the upgrade algorithm.

The final step in this process is to connect and activate any new external ports that result from the addition of new ports in this upgrade process.

The upgrade process is refined in Fig. 7B with flowchart in Fig. 12. The first part with depicting the port addition and the last step involving the activation of new external ports are the same as the more general algorithm. The distinction is that the rewiring and the optional relabeling is made more systematic. The first refinement is that the process takes place row by row or more precisely interconnection network by interconnection network. This is illustrated by the creation of the "row_select" component. Once a row is selected the optional "relabel_ports" component may be executed for that row. After that, a port that is part of the selected interconnection network, by convention the interconnection network below the selected row, is selected by the "select_port" component and rewired to its corresponding port. These three component algorithms are given in more detail below.

In the examples given below, the process of "diverting traffic" away from a port and "stop diverting traffic" from a port (also called "allow traffic to flow" to a port as stated in Fig. 9C.) should occur automatically whenever a connection is broken (disconnected) and when a connection is established, respectively if the routers are running standard routing protocols such as OSPF. However, the disconnection process may be streamlined by forcing new routing table information to be the routers prior to the connection being broken, thereby speeding the upgrade process and minimizing the impact on traffic. Likewise, a forced change to the routing table may expedite the full use of bandwidth after a new connection is established. In addition, to streamline the notation in the flowcharts, Fig. 14 and Fig. 15, the act of disconnection also implies the diversion of traffic from its attached ports prior to the breaking of the connection. Also, the act of connecting two ports implies the allowing of new traffic flow through the two ports (or equivalently the "stopping of the diversion of traffic" from the two ports) after the attachment of a connection between the two ports.

The "row_select" algorithms selects a sequencing of the interconnections between rows. Each interconnection shall be called by the row number, whose bottom ports the interconnections are attached to. This sequence must include all the interconnections, but may be in any order. Fig. 8A shows the best "row_select" algorithm. Fig. 8B shows another common "row_select" algorithm. Fig. 13A and Fig. 13B show the respective flowcharts.

There is an optional relabel_ports stage, which will be described later. The remaining component of the "upgrade" algorithm is the "select_port" port algorithm. During the rewiring of each interconnection network, the basic "upgrade" algorithm calls for the selection of any port, top or bottom not properly connected to it's appropriate corresponding port in accordance with 20 of Fig. 3B. Though any port may be selected at each step, it has been found that priority should be given to ports that are not already connected. If a port that is already connected is selected, it may introduce an additional disconnection resulting in degradation of throughput bandwidth. Fig. 9A, Fig. 9B and Fig. 9C show three common "select_port" algorithms. Corresponding flowcharts are shown in Fig. 14A, Fig. 14B, and Fig. 14C.

Though there are many possible combinations of upgrading, even based on the two "row_select" algorithms, and three "select_port" algorithms already described here. The best combination is described in detail, employing the algorithm in Fig. 8A for the row selection and the algorithm in

Fig. 9A for port selection. Since all three interconnections, 10, 11, and 12 of Fig. 3A are identical, the rewiring is shown only in detail for the first iteration of rewiring.

The process begins by selecting bottom port 3 of R(1,4) because scanning from right to left, it is the first port bottom or top not connected to anything and whose corresponding port (according to Fig. 3B) is also not connected to anything; establishing a new connection between bottom port 3 of R(1,4) and top port 3 of R(2,4) as shown in Fig. 4A; starting bottom port 3 of R(1,4) and top port 3 of R(2,4); and stop diverting the traffic from bottom port 3 of R(1,4) and top port 3 of R(2,4).

The process continues by selecting bottom port 3 of R(1,3) because scanning from right to left, it is the first port bottom or top not connected to anything and whose corresponding port (according to Fig. 3B) is also not connected to anything; establishing a new connection between bottom port 3 of R(1,3) and top port 3 of R(2,0) as shown in Fig. 4B; starting bottom port 3 of R(1,3) and top port 3 of R(2,0); and stop diverting the traffic from bottom port 3 of R(1,3) and top port 3 of R(2,0).

Scanning from right to left, no port, top or bottom, has a corresponding port which is not connected. Again, scanning from right to left, top port 3 of R(2,3) is not connected to anything with corresponding port according to Fig. 3B, bottom port 2 of R(1,4). Bottom port 2 of R(1,4) is currently connected to top port 2 of R(2,4). The breaking of this connection does not leave router R(1,4) or router R(2,4) with more than one broken connection. This completes the selection process for this step.

With top port 3 of R(2,3) selected and recalling that its corresponding port, bottom port 2 of R(1,4), is connected to top port 2 of R(2,4), the process continues by diverting traffic from bottom port 2 of R(1,4) and top port 2 of R(2,4); stopping bottom port 2 of R(1,4) and top port 2 of R(2,4); disconnecting bottom port 2 of R(1,4) and top port 2 of R(2,4) and moving the disconnected connection to top port 3 of R(2,3) as shown in Fig. 4C; starting top port 3 of R(2,3) and bottom port 2 of R(1,4); and stop diverting the traffic from top port 3 of R(2,3) and bottom port 2 of R(1,4).

Scanning from right to left, no port, top or bottom, has a corresponding port which is not connected. Again, scanning from right to left, top port 2 of R(2,4) is not connected to anything with corresponding port according to Fig. 3B, bottom port 2 of R(1,3). Bottom port 2 of R(1,3) is currently connected to top port 2 of R(2,1). The breaking this connection introduces a second broken connection to R(2,1), so top port 2 of R(2,4) is not selected. Continuing to scan from right to left, bottom port 3 of R(1,2) is not connected to anything with corresponding port according to Fig. 3B, top port 2 of R(2,1). Top port 2 of R(2,1) is currently connected to bottom port 2 of R(1,3). The breaking this connection introduces a second broken connection to R(2,1), so bottom port 3 of R(1,2) is not selected. Continuing to scan from right to left, top port 3 of R(2,2) is not connected to anything with corresponding port according to Fig. 3B, bottom port 1 of R(1,4). Bottom port 1 of R(1,4) is currently connected to top port 2 of R(2,3). The breaking of this connection does not leave router R(1,4) or router R(2,3) with more than one broken connection. This completes the selection process for this step.

With top port 3 of R(2,2) selected and recalling that its corresponding port, bottom port 1 of R(1,4), is connected to top port 2 of R(2,3), the process continues by diverting traffic from

bottom port 1 of R(1,4) and top port 2 of R(2,3); stopping bottom port 1 of R(1,4) and top port 2 of R(2,3); disconnecting bottom port 1 of R(1,4) and top port 2 of R(2,3) and moving the disconnected connection to top port 3 of R(2,2) as shown in Fig. 4D; starting top port 3 of R(2,2) and bottom port 1 of R(1,4); and stop diverting the traffic from top port 3 of R(2,2) and bottom port 1 of R(1,4).

Scanning from right to left, no port, top or bottom, has a corresponding port which is not connected. Again, scanning from right to left, top port 2 of R(2,4) is not connected to anything with corresponding port according to Fig. 3B, bottom port 2 of R(1,3). Bottom port 2 of R(1,3) is currently connected to top port 2 of R(2,1). The breaking this connection introduces a second broken connection to R(2,1), so top port 2 of R(2,4) is not selected. Continuing to scan from right to left, top port 2 of R(2,3) is not connected to anything with corresponding port according to Fig. 3B, bottom port 1 of R(1,3). Bottom port 1 of R(1,3) is currently connected to top port 2 of R(2,0). The breaking of this connection does not leave router R(1,3) or router R(2,0) with more than one broken connection. This completes the selection process for this step.

With top port 2 of R(2,3) selected and recalling that its corresponding port, bottom port 1 of R(1,3), is connected to top port 2 of R(2,0), the process continues by diverting traffic from bottom port 1 of R(1,3) and top port 2 of R(2,0); stopping bottom port 1 of R(1,3) and top port 2 of R(2,0); disconnecting bottom port 1 of R(1,3) and top port 2 of R(2,0) and moving the disconnected connection to top port 2 of R(2,3) as shown in Fig. 4E; starting top port 2 of R(2,3) and bottom port 1 of R(1,3); and stop diverting the traffic from top port 2 of R(2,3) and bottom port 1 of R(1,3).

Scanning from right to left, no port, top or bottom, has a corresponding port which is not connected. Again, scanning from right to left, top port 2 of R(2,4) is not connected to anything with corresponding port according to Fig. 3B, bottom port 2 of R(1,3). Bottom port 2 of R(1,3) is currently connected to top port 2 of R(2,1). The breaking this connection introduces a second broken connection to R(2,1), so top port 2 of R(2,4) is not selected. Continuing to scan from right to left, bottom port 3 of R(1,2) is not connected to anything with corresponding port according to Fig. 3B, top port 2 of R(2,1). Top port 2 of R(2,1) is currently connected to bottom port 2 of R(1,3). The breaking this connection introduces a second broken connection to R(2,1), so bottom port 3 of R(1,2) is not selected. Continuing to scan from right to left, bottom port 3 of R(1,1) is not connected to anything with corresponding port according to Fig. 3B, top port 1 of R(2,2). Top port 1 of R(2,2) is currently connected to bottom port 1 of R(1,2). The breaking this connection introduces a second broken connection to R(1,2), so bottom port 3 of R(1,1) is not selected. Continuing to scan from right to left, top port 3 of R(2,1) is not connected to anything with corresponding port according to Fig. 3B, bottom port 0 of R(1,4). Bottom port 0 of R(1,4) is currently connected to top port 2 of R(2,2). The breaking of this connection does not leave router R(1,4) or router R(2,2) with more than one broken connection. This completes the selection process for this step.

With top port 3 of R(2,1) selected and recalling that its corresponding port, bottom port 0 of R(1,4), is connected to top port 2 of R(2,2), the process continues by diverting traffic from bottom port 0 of R(1,4) and top port 2 of R(2,2); stopping bottom port 0 of R(1,4) and top port 2 of R(2,2); disconnecting bottom port 0 of R(1,4) and top port 2 of R(2,2) and moving the disconnected connection to top port 3 of R(2,1) as shown in Fig. 4F; starting top port 3 of R(2,1)

and bottom port 0 of R(1,4); and stop diverting the traffic from top port 3 of R(2,1) and bottom port 0 of R(1,4).

Scanning from right to left, no port, top or bottom, has a corresponding port which is not connected. Again, scanning from right to left, top port 2 of R(2,4) is not connected to anything with corresponding port according to Fig. 3B, bottom port 2 of R(1,3). Bottom port 2 of R(1,3) is currently connected to top port 2 of R(2,1). The breaking of this connection does not leave router R(1,3) or router R(2,1) with more than one broken connection. This completes the selection process for this step.

With top port 2 of R(2,4) selected and recalling that its corresponding port, bottom port 2 of R(1,3), is connected to top port 2 of R(2,1), the process continues by diverting traffic from bottom port 2 of R(1,3) and top port 2 of R(2,1); stopping bottom port 2 of R(1,3) and top port 2 of R(2,1); disconnecting bottom port 2 of R(1,3) and top port 2 of R(2,1) and moving the disconnected connection to top port 2 of R(2,4) as shown in Fig. 4G; starting top port 2 of R(2,4) and bottom port 2 of R(1,3); and stop diverting the traffic from top port 2 of R(2,4) and bottom port 2 of R(1,3).

The process continues by selecting bottom port 3 of R(1,2) because scanning from right to left, it is the first port bottom or top not connected to anything and whose corresponding port (according to Fig. 3B) is also not connected to anything; establishing a new connection between bottom port 3 of R(1,2) and top port 2 of R(2,1) as shown in Fig. 4H; starting bottom port 3 of R(1,2) and top port 2 of R(2,1); and stop diverting the traffic from bottom port 3 of R(1,2) and top port 2 of R(2,1).

Scanning from right to left, no port, top or bottom, has a corresponding port which is not connected. Again, scanning from right to left, top port 2 of R(2,2) is not connected to anything with corresponding port according to Fig. 3B, bottom port 0 of R(1,3). Bottom port 0 of R(1,3) is currently connected to top port 1 of R(2,4). The breaking of this connection does not leave router R(1,3) or router R(2,4) with more than one broken connection. This completes the selection process for this step.

With top port 2 of R(2,2) selected and recalling that its corresponding port, bottom port 0 of R(1,3), is connected to top port 1 of R(2,4), the process continues by diverting traffic from bottom port 0 of R(1,3) and top port 1 of R(2,4); stopping bottom port 0 of R(1,3) and top port 1 of R(2,4); disconnecting bottom port 0 of R(1,3) and top port 1 of R(2,4) and moving the disconnected connection to top port 2 of R(2,2) as shown in Fig. 4I; starting top port 2 of R(2,2) and bottom port 0 of R(1,3); and stop diverting the traffic from top port 2 of R(2,2) and bottom port 0 of R(1,3).

Scanning from right to left, no port, top or bottom, has a corresponding port which is not connected. Again, scanning from right to left, top port 1 of R(2,4) is not connected to anything with corresponding port according to Fig. 3B, bottom port 1 of R(1,2). Bottom port 1 of R(1,2) is currently connected to top port 1 of R(2,2). The breaking of this connection does not leave router R(1,2) or router R(2,2) with more than one broken connection. This completes the selection process for this step.

With top port 1 of R(2,4) selected and recalling that its corresponding port, bottom port 1 of R(1,2), is connected to top port 1 of R(2,2), the process continues by diverting traffic from bottom port 1 of R(1,2) and top port 1 of R(2,2); stopping bottom port 1 of R(1,2) and top port 1 of R(2,2); disconnecting bottom port 1 of R(1,2) and top port 1 of R(2,2) and moving the disconnected connection to top port 1 of R(2,4) as shown in Fig. 4J; starting top port 1 of R(2,4) and bottom port 1 of R(1,2); and stop diverting the traffic from top port 1 of R(2,4) and bottom port 1 of R(1,2).

The process continues by selecting top port 1 of R(2,2) because scanning from right to left, it is the first port bottom or top not connected to anything and whose corresponding port (according to Fig. 3B) is also not connected to anything; establishing a new connection between top port 1 of R(2,2) and bottom port 3 of R(1,1) as shown in Fig. 4K; starting top port 1 of R(2,2) and bottom port 3 of R(1,1); and stop diverting the traffic from top port 1 of R(2,2) and bottom port 3 of R(1,1).

Scanning from right to left, no port, top or bottom, has a corresponding port which is not connected. Again, scanning from right to left, bottom port 3 of R(1,0) is not connected to anything with corresponding port according to Fig. 3B, top port 0 of R(2,3). Top port 0 of R(2,3) is currently connected to bottom port 0 of R(1,1). The breaking of this connection does not leave router R(1,1) or router R(2,3) with more than one broken connection. This completes the selection process for this step.

With bottom port 3 of R(1,0) selected and recalling that its corresponding port, top port 0 of R(2,3), is connected to bottom port 0 of R(1,1), the process continues by diverting traffic from top port 0 of R(2,3) and bottom port 0 of R(1,1); stopping top port 0 of R(2,3) and bottom port 0 of R(1,1); disconnecting top port 0 of R(2,3) and bottom port 0 of R(1,1) and moving the disconnected connection to bottom port 3 of R(1,0) as shown in Fig. 4L; starting bottom port 3 of R(1,0) and top port 0 of R(2,3); and stop diverting the traffic from bottom port 3 of R(1,0) and top port 0 of R(2,3).

Scanning from right to left, no port, top or bottom, has a corresponding port which is not connected. Again, scanning from right to left, bottom port 0 of R(1,1) is not connected to anything with corresponding port according to Fig. 3B, top port 0 of R(2,4). Top port 0 of R(2,4) is currently connected to bottom port 1 of R(1,1). The breaking this connection introduces a second broken connection to R(1,1), so bottom port 0 of R(1,1) is not selected. Continuing to scan from right to left, top port 2 of R(2,0) is not connected to anything with corresponding port according to Fig. 3B, bottom port 2 of R(1,2). Bottom port 2 of R(1,2) is currently connected to top port 1 of R(2,3). The breaking of this connection does not leave router R(1,2) or router R(2,3) with more than one broken connection. This completes the selection process for this step.

With top port 2 of R(2,0) selected and recalling that its corresponding port, bottom port 2 of R(1,2), is connected to top port 1 of R(2,3), the process continues by diverting traffic from bottom port 2 of R(1,2) and top port 1 of R(2,3); stopping bottom port 2 of R(1,2) and top port 1 of R(2,3); disconnecting bottom port 2 of R(1,2) and top port 1 of R(2,3) and moving the disconnected connection to top port 2 of R(2,0) as shown in Fig. 4M; starting top port 2 of R(2,0) and bottom port 2 of R(1,2); and stop diverting the traffic from top port 2 of R(2,0) and bottom port 2 of R(1,2).

Scanning from right to left, no port, top or bottom, has a corresponding port which is not connected. Again, scanning from right to left, top port 1 of R(2,3) is not connected to anything with corresponding port according to Fig. 3B, bottom port 0 of R(1,2). Bottom port 0 of R(1,2) is currently connected to top port 1 of R(2,1). The breaking of this connection does not leave router R(1,2) or router R(2,1) with more than one broken connection. This completes the selection process for this step.

With top port 1 of R(2,3) selected and recalling that its corresponding port, bottom port 0 of R(1,2), is connected to top port 1 of R(2,1), the process continues by diverting traffic from bottom port 0 of R(1,2) and top port 1 of R(2,1); stopping bottom port 0 of R(1,2) and top port 1 of R(2,1); disconnecting bottom port 0 of R(1,2) and top port 1 of R(2,1) and moving the disconnected connection to top port 1 of R(2,3) as shown in Fig. 4N; starting top port 1 of R(2,3) and bottom port 0 of R(1,2); and stop diverting the traffic from top port 1 of R(2,3) and bottom port 0 of R(1,2).

Scanning from right to left, no port, top or bottom, has a corresponding port which is not connected. Again, scanning from right to left, top port 1 of R(2,1) is not connected to anything with corresponding port according to Fig. 3B, bottom port 2 of R(1,1). Bottom port 2 of R(1,1) is currently connected to top port 1 of R(2,0). The breaking this connection introduces a second broken connection to R(1,1), so top port 1 of R(2,1) is not selected. Continuing to scan from right to left, bottom port 0 of R(1,1) is not connected to anything with corresponding port according to Fig. 3B, top port 0 of R(2,4). Top port 0 of R(2,4) is currently connected to bottom port 1 of R(1,1). The breaking this connection introduces a second broken connection to R(1,1), so bottom port 0 of R(1,1) is not selected. Since none of the ports that are not connected satisfy the second port selection criterion, the selection process continues by selecting the rightmost port not connected to anything, that is top port 1 of R(2,1).

With top port 1 of R(2,1) selected and recalling that its corresponding port, bottom port 2 of R(1,1), is connected to top port 1 of R(2,0), the process continues by diverting traffic from bottom port 2 of R(1,1) and top port 1 of R(2,0); stopping bottom port 2 of R(1,1) and top port 1 of R(2,0); disconnecting bottom port 2 of R(1,1) and top port 1 of R(2,0) and moving the disconnected connection to top port 1 of R(2,1) as shown in Fig. 4O; starting top port 1 of R(2,1) and bottom port 2 of R(1,1); and stop diverting the traffic from top port 1 of R(2,1) and bottom port 2 of R(1,1).

Scanning from right to left, no port, top or bottom, has a corresponding port which is not connected. Again, scanning from right to left, bottom port 0 of R(1,1) is not connected to anything with corresponding port according to Fig. 3B, top port 0 of R(2,4). Top port 0 of R(2,4) is currently connected to bottom port 1 of R(1,1). The breaking this connection introduces a second broken connection to R(1,1), so bottom port 0 of R(1,1) is not selected. Continuing to scan from right to left, top port 1 of R(2,0) is not connected to anything with corresponding port according to Fig. 3B, bottom port 1 of R(1,1). Bottom port 1 of R(1,1) is currently connected to top port 0 of R(2,4). The breaking this connection introduces a second broken connection to R(1,1), so top port 1 of R(2,0) is not selected. Since none of the ports that are not connected satisfy the second port selection criterion, the selection process continues by selecting the rightmost port not connected to anything, that is bottom port 0 of R(1,1).

With bottom port 0 of R(1,1) selected and recalling that its corresponding port, top port 0 of R(2,4), is connected to bottom port 1 of R(1,1), the process continues by diverting traffic from top port 0 of R(2,4) and bottom port 1 of R(1,1); stopping top port 0 of R(2,4) and bottom port 1 of R(1,1); disconnecting top port 0 of R(2,4) and bottom port 1 of R(1,1) and moving the disconnected connection to bottom port 0 of R(1,1) as shown in Fig. 4P; starting bottom port 0 of R(1,1) and top port 0 of R(2,4); and stop diverting the traffic from bottom port 0 of R(1,1) and top port 0 of R(2,4).

The process continues by selecting bottom port 1 of R(1,1) because scanning from right to left, it is the first port bottom or top not connected to anything and whose corresponding port (according to Fig. 3B) is also not connected to anything; establishing a new connection between bottom port 1 of R(1,1) and top port 1 of R(2,0) as shown in Fig. 4Q; starting bottom port 1 of R(1,1) and top port 1 of R(2,0); and stop diverting the traffic from bottom port 1 of R(1,1) and top port 1 of R(2,0). This completes the rewiring of interconnection network 11, between row 1 and row 2.

According to the "row select" algorithm shown in Fig. 9A, interconnection network 10 should be rewired next. This is done by applying detailed rewiring just illustrated for interconnection network 11 to interconnection network 10. The results are shown in Fig. 5A.

Continuing, according to the "row select" algorithm shown in Fig. 9A, interconnection network 12 should be rewired next. This is done by applying detailed rewiring just illustrated for interconnection network 11 to interconnection network 12. The results are shown in Fig. 5B.

This completes rewiring phase. All that is needed to complete the fanout upgrade of the network shown in Fig. 3A is to connect and activate the external ports 13 and 14.

In a more sophisticated system, connections to ports on the same routers may be exchanged logically rather than physically. An example of such a swap is illustrated when interconnection network 11 of Fig. 3A is transformed into the interconnection network shown in Fig. 6A by swapping bottom port 0 and bottom port 1 of router R(1,1). This may be accomplished logically by permanently diverting the traffic intended for bottom port 0 of router R(1,1) to bottom port 1 of router R(1,1) and the traffic intended for bottom port 1 of router R(1,1) to bottom port 0 of router R(1,1). Even if such a logical exchange is not possible the ports may be in such physical proximity to each other that a physical swap can be made while minimizing the time which the disconnections occur resulting in very little impact on throughput bandwidth. In addition, only port exchanges are used here, but natural extensions to general port permutations may also be employed.

Again, rather than show a complete upgrade employing the relabel, for the entire network. The upgrade and rewiring of interconnection network 11 of Fig. 3A (which is the same as interconnection networks 10, and 12 of Fig. 3A.) is shown in great detail here. The detail algorithm is given in Fig. 10 and flow chart Fig. 15. Though the relabeling algorithm shown here starts by scanning bottom ports and scanning them from left to right. Neither is necessary, top ports can be considered first and the order of scanning just needs to be systematic, a linear scan is the simplest.

The process begins by scanning bottom ports from left to right, R(1,1) is connected to R(2,4), but according to Fig. 3B, bottom port 0 should be connected to R(2,4) instead of bottom port 1. The connection to those ports are swapped as shown in Fig. 6A.

5 The process continues by scanning bottom ports from left to right, R(1,1) is connected to R(2,0), but according to Fig. 3B, bottom port 1 should be connected to R(2,0) instead of bottom port 2. The connection to those ports are swapped as shown in Fig. 6B.

10 The process continues by scanning bottom ports from left to right, R(1,2) is connected to R(2,1), but according to Fig. 3B, bottom port 3 should be connected to R(2,1) instead of bottom port 0. The connection to those ports are swapped as shown in Fig. 6C.

15 The process continues by scanning bottom ports from left to right, R(1,2) is connected to R(2,3), but according to Fig. 3B, bottom port 0 should be connected to R(2,3) instead of bottom port 2. The connection to those ports are swapped as shown in Fig. 6D.

20 The process continues by scanning bottom ports from left to right, R(1,3) is connected to R(2,4), but according to Fig. 3B, bottom port 2 should be connected to R(2,4) instead of bottom port 0. The connection to those ports are swapped as shown in Fig. 6E.

25 The process continues by scanning bottom ports from left to right, R(1,3) is connected to R(2,0), but according to Fig. 3B, bottom port 3 should be connected to R(2,0) instead of bottom port 1. The connection to those ports are swapped as shown in Fig. 6F.

30 The process continues by scanning bottom ports from left to right, R(1,4) is connected to R(2,2), but according to Fig. 3B, bottom port 1 should be connected to R(2,2) instead of bottom port 0. The connection to those ports are swapped as shown in Fig. 6G.

35 The process continues by scanning bottom ports from left to right, R(1,4) is connected to R(2,3), but according to Fig. 3B, bottom port 2 should be connected to R(2,3) instead of bottom port 0. The connection to those ports are swapped as shown in Fig. 6H.

40 The process continues by scanning bottom ports from left to right, R(1,4) is connected to R(2,4), but according to Fig. 3B, bottom port 3 should be connected to R(2,4) instead of bottom port 0. The connection to those ports are swapped as shown in Fig. 6I.

45 Scanning from left to right all bottom ports that should be swapped are swapped. In this example, the next step occurs by scanning top ports now from right to left. The direction is not important, but is reversed here purely for aesthetic value. In fact, any systematic scan will work. Unlike the scanning of bottom ports, the top ports should be connected to the correct bottom port, not just the correct router.

The process continues by scanning top ports from right to left, R(2,4) is connected to bottom port 3 of R(1,4), but according to Fig. 3B, top port 3 should be connected to bottom port 3 of R(1,4) instead of top port 2. The connection to those ports are swapped as shown in Fig. 6J.

The process continues by scanning top ports from right to left, R(2,4) is connected to bottom port 2 of R(1,3), but according to Fig. 3B, top port 2 should be connected to bottom port 2 of R(1,3) instead of top port 1. The connection to those ports are swapped as shown in Fig. 6K.

- 5 The process continues by scanning top ports from right to left, R(2,3) is connected to bottom port 2 of R(1,4), but according to Fig. 3B, top port 3 should be connected to bottom port 2 of R(1,4) instead of top port 2. The connection to those ports are swapped as shown in Fig. 6L.

- 10 The process continues by scanning top ports from right to left, R(2,2) is connected to bottom port 1 of R(1,4), but according to Fig. 3B, top port 3 should be connected to bottom port 1 of R(1,4) instead of top port 2. The connection to those ports are swapped as shown in Fig. 6M.

- 15 The process continues by scanning top ports from right to left, R(2,1) is connected to bottom port 3 of R(1,2), but according to Fig. 3B, top port 2 should be connected to bottom port 3 of R(1,2) instead of top port 1. The connection to those ports are swapped as shown in Fig. 6N.

The process continues by scanning top ports from right to left, R(2,0) is connected to bottom port 3 of R(1,3), but according to Fig. 3B, top port 3 should be connected to bottom port 3 of R(1,3) instead of top port 2. The connection to those ports are swapped as shown in Fig. 6O.

This completes the scanning of the top ports, and the relabeling process. The process continues by employing the same "select_port" algorithm as in Fig. 9A as before.

The process continues by selecting bottom port 0 of R(1,4) because scanning from right to left, it is the first port bottom or top not connected to anything and whose corresponding port (according to Fig. 3B) is also not connected to anything; establishing a new connection between bottom port 0 of R(1,4) and top port 3 of R(2,1) as shown in Fig. 6P; starting bottom port 0 of R(1,4) and top port 3 of R(2,1); and stop diverting the traffic from bottom port 0 of R(1,4) and top port 3 of R(2,1).

- 35 The process continues by selecting top port 2 of R(2,3) because scanning from right to left, it is the first port bottom or top not connected to anything and whose corresponding port (according to Fig. 3B) is also not connected to anything; establishing a new connection between top port 2 of R(2,3) and bottom port 1 of R(1,3) as shown in Fig. 6Q; starting top port 2 of R(2,3) and bottom port 1 of R(1,3); and stop diverting the traffic from top port 2 of R(2,3) and bottom port 1 of R(1,3).

- 40 The process continues by selecting bottom port 2 of R(1,2) because scanning from right to left, it is the first port bottom or top not connected to anything and whose corresponding port (according to Fig. 3B) is also not connected to anything; establishing a new connection between bottom port 2 of R(1,2) and top port 2 of R(2,0) as shown in Fig. 6R; starting bottom port 2 of R(1,2) and top port 2 of R(2,0); and stop diverting the traffic from bottom port 2 of R(1,2) and top port 2 of R(2,0).

- 45 Scanning from right to left, no port, top or bottom, has a corresponding port which is not connected. Again, scanning from right to left, top port 1 of R(2,4) is not connected to anything with corresponding port according to Fig. 3B, bottom port 1 of R(1,2). Bottom port 1 of R(1,2) is currently connected to top port 1 of R(2,2). The breaking this connection introduces a second

broken connection to R(2,2), so top port 1 of R(2,4) is not selected. Continuing to scan from right to left, top port 2 of R(2,2) is not connected to anything with corresponding port according to Fig. 3B, bottom port 0 of R(1,3). Bottom port 0 of R(1,3) is currently connected to top port 1 of R(2,1). The breaking of this connection does not leave router R(1,3) or router R(2,1) with more than one broken connection. This completes the selection process for this step.

With top port 2 of R(2,2) selected and recalling that its corresponding port, bottom port 0 of R(1,3), is connected to top port 1 of R(2,1), the process continues by diverting traffic from bottom port 0 of R(1,3) and top port 1 of R(2,1); stopping bottom port 0 of R(1,3) and top port 1 of R(2,1); disconnecting bottom port 0 of R(1,3) and top port 1 of R(2,1) and moving the disconnected connection to top port 2 of R(2,2) as shown in Fig. 6S; starting top port 2 of R(2,2) and bottom port 0 of R(1,3); and stop diverting the traffic from top port 2 of R(2,2) and bottom port 0 of R(1,3).

Scanning from right to left, no port, top or bottom, has a corresponding port which is not connected. Again, scanning from right to left, top port 1 of R(2,4) is not connected to anything with corresponding port according to Fig. 3B, bottom port 1 of R(1,2). Bottom port 1 of R(1,2) is currently connected to top port 1 of R(2,2). The breaking of this connection does not leave router R(1,2) or router R(2,2) with more than one broken connection. This completes the selection process for this step.

With top port 1 of R(2,4) selected and recalling that its corresponding port, bottom port 1 of R(1,2), is connected to top port 1 of R(2,2), the process continues by diverting traffic from bottom port 1 of R(1,2) and top port 1 of R(2,2); stopping bottom port 1 of R(1,2) and top port 1 of R(2,2); disconnecting bottom port 1 of R(1,2) and top port 1 of R(2,2) and moving the disconnected connection to top port 1 of R(2,4) as shown in Fig. 6T; starting top port 1 of R(2,4) and bottom port 1 of R(1,2); and stop diverting the traffic from top port 1 of R(2,4) and bottom port 1 of R(1,2).

The process continues by selecting top port 1 of R(2,2) because scanning from right to left, it is the first port bottom or top not connected to anything and whose corresponding port (according to Fig. 3B) is also not connected to anything; establishing a new connection between top port 1 of R(2,2) and bottom port 3 of R(1,1) as shown in Fig. 6U; starting top port 1 of R(2,2) and bottom port 3 of R(1,1); and stop diverting the traffic from top port 1 of R(2,2) and bottom port 3 of R(1,1).

Scanning from right to left, no port, top or bottom, has a corresponding port which is not connected. Again, scanning from right to left, top port 1 of R(2,1) is not connected to anything with corresponding port according to Fig. 3B, bottom port 2 of R(1,1). Bottom port 2 of R(1,1) is currently connected to top port 0 of R(2,3). The breaking of this connection does not leave router R(1,1) or router R(2,3) with more than one broken connection. This completes the selection process for this step.

With top port 1 of R(2,1) selected and recalling that its corresponding port, bottom port 2 of R(1,1), is connected to top port 0 of R(2,3), the process continues by diverting traffic from bottom port 2 of R(1,1) and top port 0 of R(2,3); stopping bottom port 2 of R(1,1) and top port 0 of R(2,3); disconnecting bottom port 2 of R(1,1) and top port 0 of R(2,3) and moving the disconnected connection to top port 1 of R(2,1) as shown in Fig. 6V; starting top port 1 of R(2,1)

and bottom port 2 of R(1,1); and stop diverting the traffic from top port 1 of R(2,1) and bottom port 2 of R(1,1).

The process continues by selecting top port 0 of R(2,3) because scanning from right to left, it is the first port bottom or top not connected to anything and whose corresponding port (according to Fig. 3B) is also not connected to anything; establishing a new connection between top port 0 of R(2,3) and bottom port 3 of R(1,0) as shown in Fig. 6W; starting top port 0 of R(2,3) and bottom port 3 of R(1,0); and stop diverting the traffic from top port 0 of R(2,3) and bottom port 3 of R(1,0).

This completes the relabel and rewire phase applied to interconnection network 11 of Fig. 3A. In a complete upgrade, this algorithm could then be applied to interconnection network 10 of Fig. 3A, and finally interconnection network 12 of Fig. 3B in accordance with the first "row_select" algorithm of Fig. 8. After all the interconnection networks are completed, the upgrade is finished by connecting and activating external connections 13 and 14.

The process depicted in application may be used also to reduce the fanout of routers in a network. Simply reverse the process shown here. In summary, the process would involve disconnecting external connections to ports to be removed; rewiring each interconnection network in the same manner as those described above; and finally removing the unwired ports.

Though this invention is applied to a RBCCG network composed of rows of routers of consistent number per row and fanout per router. It may be applied to fault tolerant networks composed of rows of routers of varying number per row and fanout per router where the rows are connected with interconnection networks of the same type as that used in a RBCCG network. It also applies to any multistage network which can be expanded by fanout since the action of adding ports induces fault tolerance. In addition, it is not restricted to routers and may be applied to networks composed of other types of switching elements.

Although the present invention has been described above in terms of specific embodiments, it is anticipated that alteration and modifications thereof will no doubt become apparent to those skilled in the art. It is therefore intended that the following claims be interpreted as covering all such alterations and modifications as falling within the true spirit and scope of the invention.